arXiv:1704.07741v1 [cond-mat.dis-nn] 25 Apr 2017

High storage capacity in the Hopfield model with auto-interactions - stability analysis

Jacopo Rocchi

Nonlinearity and Complexity Research Group, Aston University, Birmingham B4 7ET, United Kingdom

E-mail: j.rocchi@aston.ac.uk

David Saad

Nonlinearity and Complexity Research Group, Aston University, Birmingham B4 7ET, United Kingdom

E-mail: d.saad@aston.ac.uk

Daniele Tantari

Scuola Normale Superiore, Centro Ennio de Giorgi, Piazza dei Cavalieri 3, I-56100, Pisa, Italy

E-mail: daniele.tantari@sns.it

Abstract. Recent studies point to the potential storage of a large number of patterns in the celebrated Hopfield associative memory model, well beyond the limits obtained previously. We investigate the properties of new fixed points to discover that they exhibit instabilities for small perturbations and are therefore of limited value as associative memories. Moreover, a large deviations approach also shows that errors introduced to the original patterns induce additional errors and increased corruption with respect to the stored patterns.

1. Introduction

Hopfield models [1, 2, 3] are recurrent neural networks where connections between units form a fully connected symmetric network. They have been proposed as models of content addressable memories, i.e. systems that are able to retrieve memory items from partial information. Their introduction was inspired by the observation that in large physical systems, interactions between the elementary degrees of freedom generate collective phenomena, such low temperature magnetisation in Ising models. Following this idea, the stability of memories in systems of interacting neurons has been

successfully described as an emergent property, instigated by the dynamics of neural network models [1, 2].

Any physical system whose dynamics is dominated by a number of locally stable states can act as a content addressable memory as long as these states can be controlled. The Hebbian rule [4] has played an important role in training the couplings between neurons such that a prescribed set of memories (binary configurations of the Hopfield model) become attractors of the dynamics. Hopfield pointed out [1] that the issue of pattern retrieval is non trivial and that retrieval performance falls rapidly as more patterns are introduced, and are incorporated in the couplings. This behaviour was first found in numerical simulations and then analysed utilising statistical mechanics methods and exploiting the analogy with spin glass models. For P is the number of patterns, N the number of neurons and $\alpha = P/N$, it has been found [5, 6] that the critical value α_c below which recovery is possible is approximately $\alpha_c \approx 0.14$. Original studies considered the case where diagonal terms are not present, i.e. neural network models without auto-interactions. Subsequent studies [7, 8] considered also the problem with auto-interactions, but only recently it has been pointed out [9] that in this case an interesting regime can be found at $\alpha \gg 1$. The probability that a given pattern is not a fixed point of the dynamics was studied and it has been shown that this probability is very small for very low α, as expected, but surprisingly that there is another unexplored region at very large α where this probability is very small as well. While the former and other intermediate regimes are well studied [6], the behaviour at $\alpha \gg 1$ was unexpected, since it implies that in this new regime the patterns are again fixed points of the dynamical equations. Moreover, it has been pointed out [9] that this new regime does not appear in the absence of diagonal interaction terms.

In this paper we study the stability of these fixed points. The relevance of this analysis comes from the fact that associative memories are useful when (in the regime where) they recover memories on the basis of similarity. In other words, the Hopfield model can be used to retrieve memories when starting the dynamics from a configuration $\underline{s}$ similar, but not exactly equal to, a given pattern, we converge to one of the original patterns. Our analysis suggests that although training patterns are fixed points of the dynamics in the newly discovered regime [9], they are unstable, contrarily to the known regime of small α values. In Sec. 2 we introduce the model and its dynamics; while in Sec. 3 we compute the probability of escaping the stored pattern when a small perturbation is introduced. In Sec. 4 we complete the analysis by computing the typical number of errors made after one dynamical step, where errors are measured in terms of the Hamming distance between the initial configuration (one of the training vectors) and the dynamical configuration.

2. Dynamics of a neural network

The neural network model that we will consider in this work is a system of N binary variables (neurons) $s_i \in \{-1, +1\},\ i = 1 \dots N$, interconnected by a symmetric network

of synapses specified by the real coupling matrix J_{ij}. We will focus on the non-linear dynamical equations

$$s_i(t+1) = \text{sign}(x) \left(\sum_{j=1}^{N} J_{ij} s_j(t) \right) , \tag{1}$$

where the value $s_i(t+1)$ represents the state of the neuron i at time $t+1$, which may be active, $s_i(t+1) = +1$, or inactive, $s_i(t+1) = -1$. The value $s_i(t+1)$ depends on the state of the neurons at the previous time step, $\{s_j(t)\}$. These equations give rise to a dynamical process in the space of configurations, depending on the properties of the matrix J_{ij} but, as pointed out by Hopfield [1], a careful choice of J_{ij} may trap the dynamics in basins of attraction that correspond to a given random set of patterns (training vectors) $\{\underline{\xi}^{\nu}\}_{\nu=1,\ldots,P}$ where $\xi_i^{\nu} \in [-1,+1]$, $i = 1 \ldots N$ and $\nu = 1 \ldots P$. These patterns can be considered as memories that the system is able to retrieve and should be fixed points of Eq. (1). In the following we will focus on the case where the matrix J_{ij} is specified by the Hebbian rule [4],

$$J_{ij} = \frac{1}{P} \sum_{\nu=1}^{P} \xi_i^{\nu} \xi_j^{\nu} , \tag{2}$$

introduced in order to explain *associative learning*. In fact, for $P = O(N)$, Eq. (2) can be obtained cumulatively from the successive application of the learning rule $\Delta J_{ij} \propto \xi_i^{\mu} \xi_j^{\mu}$, specifying the change in the coupling between neurons when learning a given pattern μ and describing the observation that the simultaneous activation of neurons i and j increased the coupling strength between them.

Retrieval of patterns is known to be possible only for a number of patterns that is a small fraction of that of the neurons [1, 2, 5, 6]. Diagonal interaction terms were not considered in the early works about Hopfield model for a physical reason: in the corresponding spin models the field of a variable is induced by the state of its neighbours, but not on its own; thus self interactions do not exist and $J_{ii} = 0, \forall i$. Neural networks with diagonal terms have been studied in [9] and a very interesting regime has been found for $\alpha \gg 1$. The probability p_V that a given pattern $\underline{\xi}^{\mu}$ is a not fixed point of the dynamics has been computed and has been shown to be very small for very low $\alpha = P/N$, as expected, but surprisingly another region has been identified at the very large α regime, where p_V is also very small. In the intermediate regime, p_V is large. The first and the intermediate regime are well known but the behaviour of p_V at $\alpha \gg 1$ was new and unexpected. The probability $\bar{p}_V$ that a random vector, not in the training set $\{\xi_i^{\nu}\}_{\nu=1,\ldots,P}$, is not a fixed point of the dynamics has also been studied. As expected, in the low α regime, the probability $\bar{p}_V$ is close to 1 but, and as α increases this probability vanishes. Thus, in the new, large α regime, $\bar{p}_V$ is also very small. In other words, in this regime, both patterns $\underline{\xi}^{\nu}$ and random configurations are likely to be fixed point of the dynamics. This has a trivial interpretation by noticing that for very large $P \gg N$, the interaction matrix J_{ij} defined in Eq. (2) tends to the unit matrix. Even if this result seems to invalidate the usefulness of this new regime, it was shown that the ratio $\bar{p}_V/p_V$ tends to a finite number, e, in the large N limit and that real patterns have an

higher probability of being fixed points of the model. In the next section we address the stability question of these patterns, making use of the same strategy used in [9].

3. Stability of the fixed points

In order to study the stability of fixed points of the dynamical equations given in Eq. (1), we consider the case where one of the patterns is randomly perturbed. Since the patterns $\underline{\xi}^{\nu}$ are configurations of binary variables, a random perturbation is obtained by flipping the value of K sites. We denote by $\mathcal{P}$ the set of perturbed sites and by $\mathcal{L}$ the set of unperturbed variables and clearly, $\mathcal{L} = \overline{\mathcal{P}}$. We can consider the equations

$$s_i' = \text{sign}\left(\sum_{j=1}^{N}\sum_{\nu=1}^{P}\xi_i^{\nu}\xi_j^{\nu}s_j\right) \tag{3}$$

and compute the probability to get back to ξ_i^{μ} when the starting configuration is given by

$$s_j = \xi_j^{\mu}\mathbb{I}_{j,\mathcal{L}} - \xi_j^{\mu}\mathbb{I}_{j,\mathcal{P}}\,, \tag{4}$$

where $\mathbb{I}_{j,\mathcal{P}(\mathcal{L})}$ is one if $j \in \mathcal{P}(\mathcal{L})$, and zero otherwise. As in [9] we focus on the one-step dynamics.

Let us first consider the case when $i \notin \mathcal{P}$. After some elementary algebra, the argument of the sign function in the r.h.s. of Eq. (3) becomes

$$(P+N-1-2K)\xi_i^{\mu} + \sum_{j\neq i}\sum_{\nu\neq\mu}\xi_i^{\nu}\xi_j^{\nu}\left[\xi_j^{\mu}\mathbb{I}_{j,\mathcal{L}} - \xi_j^{\mu}\mathbb{I}_{j,\mathcal{P}}\right]\,. \tag{5}$$

The second term contains $(N-1)(P-1)$ uncorrelated terms of unitary variance. Using the central limit theorem we obtain for large N

$$s_i' = \text{sign}\left(\frac{(P+N-1-2K)}{\sqrt{(N-1)(P-1)}}\xi_i^{\mu} + z\right)\,, \tag{6}$$

where $z \sim \mathcal{N}(0,1)$ is drawn from a normalised Gaussian distribution. Clearly, if z were 0, $s_i' = \xi_i^{\mu}$ and we recover the correct sign. The variable z, induced by the arbitrary bit flips, is thus a destabilising term, that impacts on the r.h.s. of Eq. (6). It is actually harmless as soon as it doesn't change the sign of ξ_i^{μ}, thus we make a mistake on the value s_i' with a probability equal to

$$p^{\mathcal{L}} = \int_{-\infty}^{-\frac{(P+N-1-2K)}{\sqrt{(N-1)(P-1)}}} P(z)dz = \frac{1}{2}\text{erfc}\left\{\frac{(P+N-1-2K)}{\sqrt{2(N-1)(P-1)}}\right\}\,. \tag{7}$$

Analogously, when $i \in \mathcal{P}$, the r.h.s. of Eq. (3) becomes

$$(N-P+1-2K)\xi_i^{\mu} + \sum_{j\neq i}\sum_{\nu\neq\mu}\xi_i^{\nu}\xi_j^{\nu}\left[\xi_j^{\mu}\mathbb{I}_{j,\mathcal{L}} - \xi_j^{\mu}\mathbb{I}_{j,\mathcal{P}}\right]\,. \tag{8}$$

Using a central limit argument one obtains

$$s_i' = \text{sign}\left(\frac{(N-P+1-2K)}{\sqrt{(N-1)(P-1)}}\xi_i^{\mu} + z\right)\,, \tag{9}$$

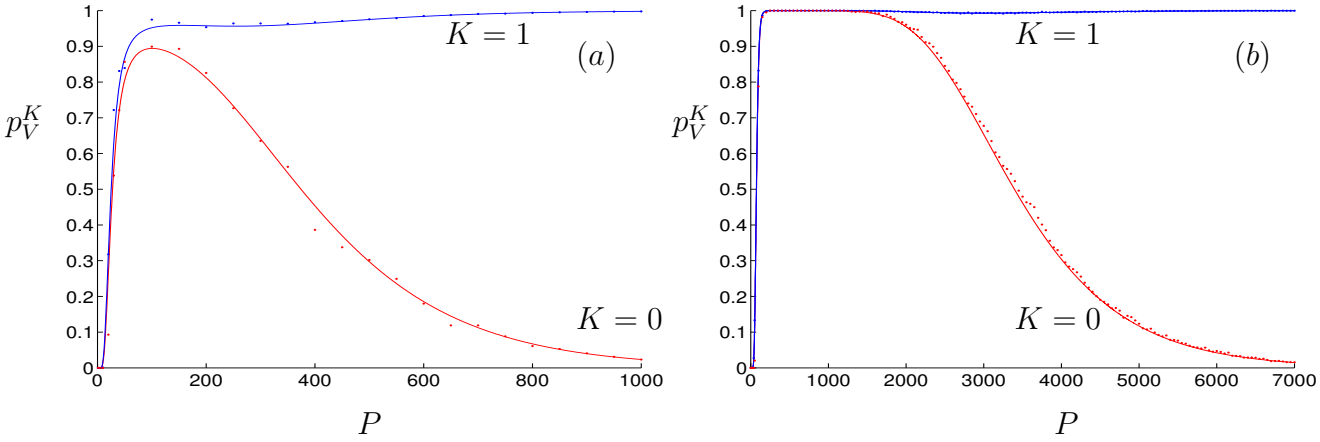

Figure 1. The probability of not recovering the stored pattern after one iteration after K variables have been perturbed p_V^K for $K = 1$ (blue line) and $K = 0$ (red line). Dots are obtained by numerical simulations, where probability p_V^K is estimated counting the number of times that Eq. (1) does not give a given pattern when we perturbed $K = 1$ (blue dots) and $K = 0$ (red dots) neurons, repeating this process 10^4 times for different patterns at any value of P. (a) For $N = 100$ and (b) For $N = 500$.

where z is again drawn from a normalised Gaussian distribution. As in Eq. (7) we obtain the probability of making an error on one of the perturbed variables,

$$p^{\mathcal{P}} = \int_{-\infty}^{-\frac{(N-P+1-2K)}{\sqrt{(N-1)(P-1)}}} P(z)dz = \frac{1}{2}\text{erfc}\left\{\frac{(N-P+1-2K)}{\sqrt{2(N-1)(P-1)}}\right\} . \tag{10}$$

Notice that $p^{\mathcal{P}}$ and $p^{\mathcal{L}}$ differ in the sign in front of P, indicating the contribution coming from the diagonal interaction components. This contribution is always aligned with the variable value and consequently decreases the error probability in unperturbed variables. In the limit of large N and $P = \alpha N$, and finite K we obtain

$$p^{\mathcal{L}} = \frac{1}{2}\text{erfc}\left(\frac{1}{\sqrt{2\alpha}} - \sqrt{\frac{\alpha}{2}}\right) \qquad p^{\mathcal{P}} = \frac{1}{2}\text{erfc}\left(\frac{1}{\sqrt{2\alpha}} + \sqrt{\frac{\alpha}{2}}\right) . \tag{11}$$

While $p^{\mathcal{L}}$ tends to zero as $\alpha \to 0$ and $\alpha \to \infty$, with a maximum at $\alpha = 1$, the error probability for perturbed spins $p^{\mathcal{P}}$ is an increasing function of α, going from 0 to 1.

This difference in stability between perturbed and unperturbed spins affects the overall stability of the original pattern. Since there are K perturbed variables and $N-K$ unperturbed variables, the probability of failing to recover the original pattern $\underline{\xi}^\mu$ after a single step of parallel dynamics is

$$p_V^K = 1 - \left(1 - p^{\mathcal{P}}\right)^K \left(1 - p^{\mathcal{L}}\right)^{N-K} \tag{12}$$

that, for $K = 0$, becomes the probability p_V [9]. This probability can be plotted for different P values at a given N. In Fig. 1(a) we plot p_V^K for $N = 100$ and $K = 1$ (blue) and $K = 0$ (red), where the second case corresponds to the unperturbed case studied in [9]. We also performed numerical simulations (dots) in systems of $N = 100$ variables for a different number of P, counting the number of times that taking one of the training vectors and perturbing $K = 1$ of its values we did not recover the original training vector, repeating this procedure 10^4 times. While at large P values the probability p_V (red line) decreases to zero, consistently with the observation made in [9], this does not happen for the perturbed case (blue line). In other words, perturbing just one variable in a system of 100 neurons is sufficient to not recover the correct patterns in the regime $P \gg N$. Moreover, we notice that for small P values both lines are close to zero, meaning that in this regime perturbing one neuron does not make a big difference. This is in agreement with Eq. (11): for $K = 0$, p_V follows the behaviour of $p^{\mathcal{L}}$, where the stability of unperturbed spins at $K > 0$ is affected by perturbed spins that are dominated by the diagonal interactions at large P values, which increase the error probability. In Fig. 1(b) we plot the same quantity p_V^K in the case $K = 0$ (red line) and $K = 1$ (blue line) for a system of $N = 500$, finding the same qualitative behaviour.

4. Large Deviations

In this section we compute the typical number of errors observed in $\underline{s}'$ with respect to the original patterns, produced by applying Eq. (3) to the vector of dynamical variables $\underline{s}$, which is specified by Eq. (4) for a given number K of perturbed spins. Let us denote by $M_{\mathcal{P}}$ the number of errors in the set of perturbed spins $\mathcal{P}$ and by $M_{\mathcal{L}}$ the number of errors in the set of unperturbed spins $\mathcal{L}$ in $\underline{s}$. While the probability of $M_{\mathcal{P}}$ is given by

$$P_{\mathcal{P}}(M_{\mathcal{P}}) = \binom{K}{M_{\mathcal{P}}} \left(p^{\mathcal{P}}\right)^{M_{\mathcal{P}}} \left(1 - p^{\mathcal{P}}\right)^{K - M_{\mathcal{P}}} , \tag{13}$$

the probability of $M_{\mathcal{L}}$ is given by

$$P_{\mathcal{L}}(M_{\mathcal{L}}) = \binom{N-K}{M_{\mathcal{L}}} \left(p^{\mathcal{L}}\right)^{M_{\mathcal{L}}} \left(1 - p^{\mathcal{L}}\right)^{N-K-M_{\mathcal{L}}} . \tag{14}$$

Let us denote the total number of error by $M = M_{\mathcal{P}} + M_{\mathcal{L}}$. The probability of M, the number of errors at the next time step, is given by

$$P(M) \propto \sum_{M_{\mathcal{P}}=0}^{K} \sum_{M_{\mathcal{L}}=0}^{N-K} P_{\mathcal{P}}(M_{\mathcal{P}}) P_{\mathcal{L}}(M_{\mathcal{L}}) \delta_{M, M_{\mathcal{P}} + M_{\mathcal{L}}} \tag{15}$$

and we readily obtain

$$P(M) \propto \sum_{M_{\mathcal{P}}=\max(0,M+K-N)}^{\min(K,M)} P_{\mathcal{P}}(M_{\mathcal{P}})P_{\mathcal{L}}(M-M_{\mathcal{P}})\,. \tag{16}$$

Since we are mostly interested in the large N behaviour, we denote by $K = N\rho$, $M = Nm$, $M_{\mathcal{P}} = Nm_{\mathcal{P}}$ and $M_{\mathcal{L}} = Nm_{\mathcal{L}}$, and use the Sterling's formula for approximating the factorial of a large integer,

$$\lim_{N\to\infty} N! = \sqrt{2\pi N}\left(\frac{N}{\mathrm{e}}\right)^N\,. \tag{17}$$

Simple algebra leads to the expression

$$P(m) \propto \int_{\max(0,m+\rho-1)}^{\min(\rho,m)} dm_{\mathcal{P}} \mathrm{e}^{N\phi(m_{\mathcal{P}},m)} \approx \mathrm{e}^{N\phi(m^*_{\mathcal{P}},m)} \tag{18}$$

where $m^*_{\mathcal{P}}$ is the maximum of $\phi(m_{\mathcal{P}},m)$ over $m_{\mathcal{P}} \in [\max(0,m+\rho-1),\min(\rho,m)]$ at a given m. The large deviation function of the probability $P(m)$ is given by $\phi(m^*_{\mathcal{P}},m) \equiv \psi(m)$ and its expression is

$$\begin{aligned}\psi(m) = {} & \rho\log(\rho) + (1-\rho)\log(1-\rho) \\ & + m^*_{\mathcal{P}}\log\left(\frac{p^{\mathcal{P}}}{m^*_{\mathcal{P}}}\right) + (\rho - m^*_{\mathcal{P}})\log\left(\frac{1-p^{\mathcal{P}}}{\rho - m^*_{\mathcal{P}}}\right) \\ & + (m - m^*_{\mathcal{P}})\log\left(\frac{p^{\mathcal{L}}}{m - m^*_{\mathcal{P}}}\right) \\ & + (1-\rho-m+m^*_{\mathcal{P}})\log\left(\frac{1-p^{\mathcal{L}}}{1-\rho-m+m^*_{\mathcal{P}}}\right)\,.\end{aligned} \tag{19}$$

Let us first consider the case $\rho = 0$, the case where the selected pattern is not perturbed at all. In Fig. 2(a) we plot

$$\omega(m) = \log(\log(-\psi(m)+1)+1) \tag{20}$$

for $N = 100$, $K = 0$ as a function of $\log(m+1)$ for different choices of P. The maximum m^* of $\psi(m)$ corresponds to a minimum of $\omega(m)$, where the double logarithm is chosen to emphasise the difference between different lines at large P values. While at small P values we find $m^* = 0$, i.e. the most likely value for m is zero, corresponding to a non-increasing number of errors, as P increases the probability of observing $m = 0$ decreases and for $P > 15$ (which corresponds roughly to $\alpha = 0.15$) we find a different minimum at $m > 0$, as can be seen in Fig. 2(b) and Fig. 3(a). We also observe that the probability of Nm is sharply peaked in 0 in the low P regime, while it is much broader in the large P regime, even if it is clearly visible that the probability of observing $Nm = 0$ is negligible.

The behaviour of m^* at $K = 10$ can be seen in Fig. 2(d) leading to a qualitatively similar behaviour of $\omega(m)$ shown in Fig. 2(c), where we observe that small values of P are dominated by $m^* = 0$, while as P increases m^* remains grater than zero. In other words, while the low P regime leads to a recovery of the original pattern with probability 1 even in the case when we perturb $K = 10$ variables, the large P regime does not. Notice

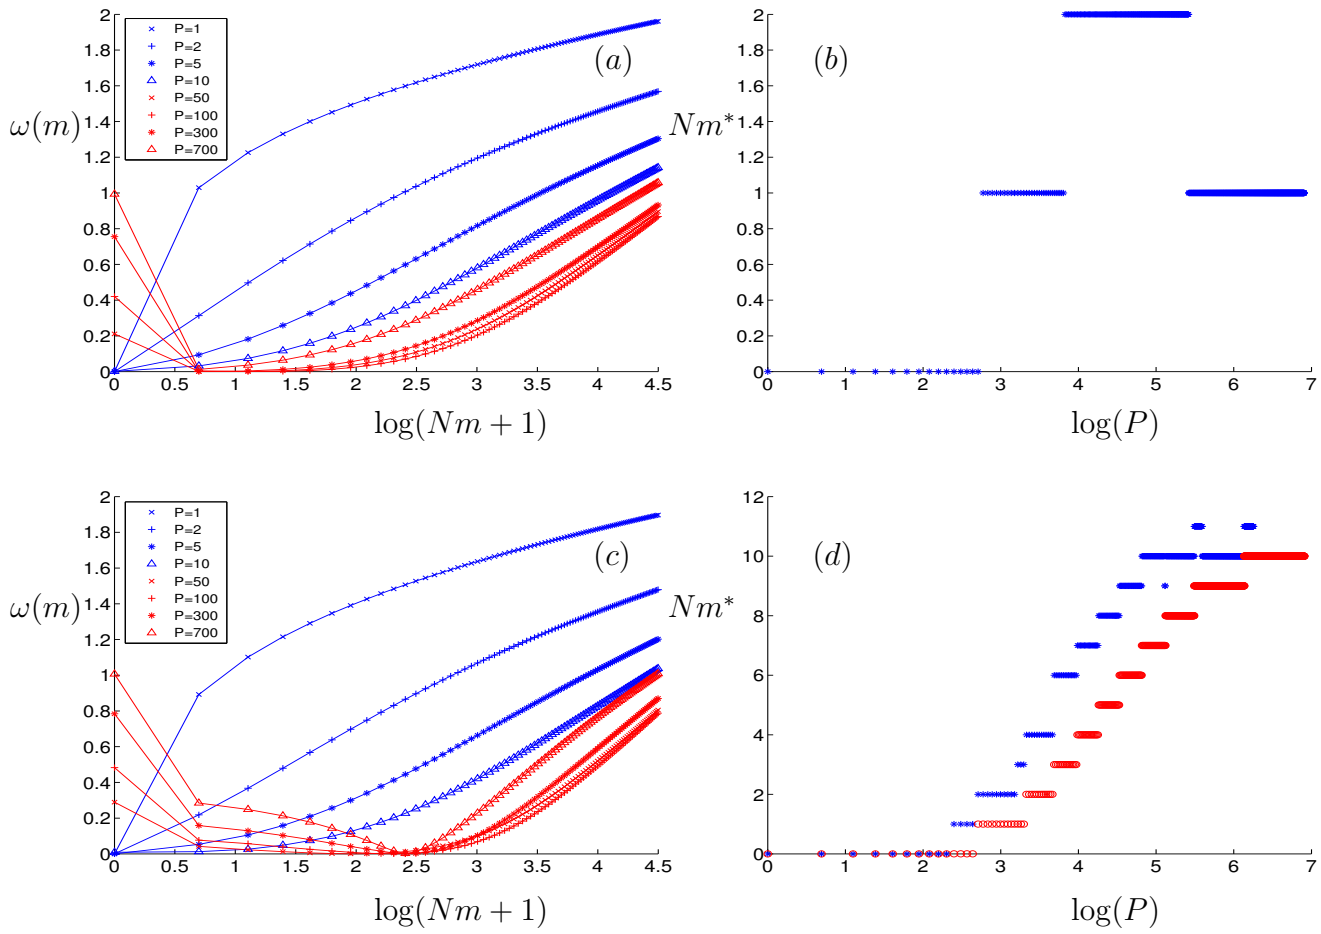

Figure 2. (a) Function $\omega(m)$, defined in Eq. (20), for different values of P at $N = 100$, $K = 0$, plotted as a function of $\log(Nm+1)$. The minimum of this function corresponds to the most likely values of Nm, i.e. the typical number of errors made after one dynamical step. (b) The dominating parameter Nm^* for $N = 100$, $K = 0$ for different (logarithmic) values of P. We see that in the large P regime $Nm^* > 0$ even when no perturbations are introduced. In this regime, the typical number of mistakes made is $Nm^* = 1$. Notice the severe flattening of the large deviation function observed in sub-figure (a) for values $Nm > 0$. (c) The same as (a) but for $K = 10$. (d) The blue x-marks indicate the dominating parameter plotted in (b) Nm^* for $K = 10$, while the red circles indicate $Nm^*_{\mathcal{P}}(m^*)$, i.e. the typical number of error made in the perturbed set of variables. For small values of P the dominating value is $m^* = 0$, implying that we correctly reduce the number of errors just after one dynamical time step and recover the original pattern with probability 1. However, in the large P regime, $Nm^* = 10$ implying that the number of errors is not reduced at all.

in Fig. (2)(d) that the value of m^* is mainly given by $m^*_{\mathcal{P}}$ with fluctuations of order $1/N$: errors are mainly induced by the set of perturbed spins that remain blocked because of the dominating diagonal terms.

Finally, to emphasise the sensitivity of fixed points to perturbations we plotted the m^* values for a single pattern error $K = 1$ shown in Fig. 3(b). We observe the same qualitative behaviour of Fig. 2(b).

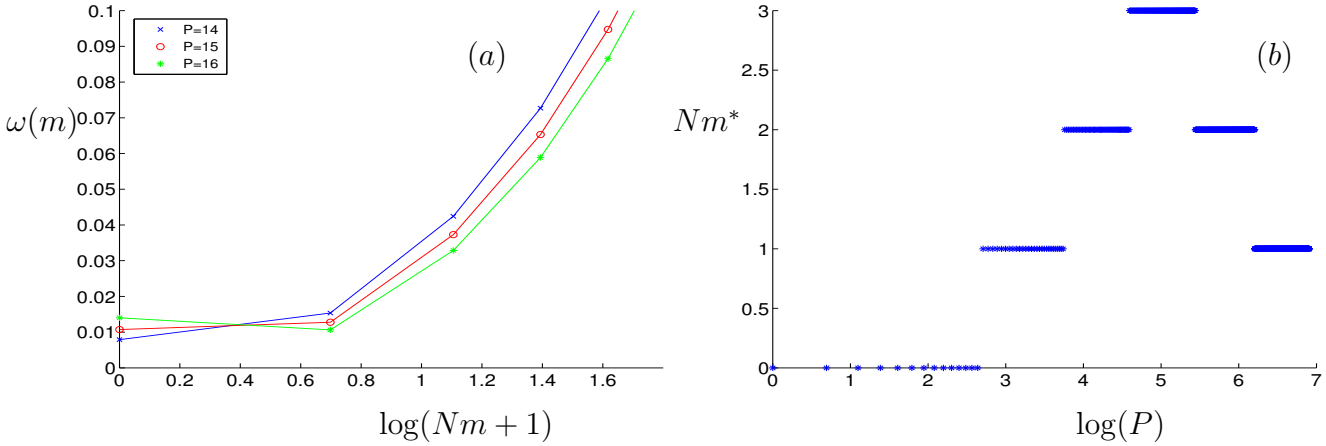

Figure 3. (a) Function $\omega(m)$, defined in Eq. (20), for different values of $P = 14, 15, 16$ at $N = 100$, $K = 0$, plotted as a function of $\log(Nm + 1)$. We see that for $P > 15$ the minimum of $\omega(m)$ becomes non-zero. (b) The dominating parameter Nm^* for $N = 100$, $K = 1$ for different values of P. The x-axis is $\log(P)$. We basically gain the same information we got from Fig. 2(b): in the large P regime the typical number of errors is greater than zero.

Conclusion

The discovery of fixed points in the Hopfield model at the large number of patterns limit raised new hopes for a high-capacity properties of the Hopfield model, especially in within the context of associative memories in neural networks. We examine the usefulness of the newly discovered fixed points by focussing on their ability to recover stored patterns on the basis of incomplete information. In other words, the ability to converge to the original pattern when starting from a configuration that is similar to, but not exactly equal to it. We study the stability properties of these fixed points and show that these fixed points are unstable with respect to small perturbations. We also investigate the typical number of errors made by the one time step dynamics given in Eq. (1) and find that while this number is zero in the low storage regimes, it is not in the new large storage regime.

Finally, we notice that a simple statistical mechanical argument suggests that it is unlikely that the phase diagrams of an Hopfield model with auto-interactions differs from the phase diagram of an Hopfield model without auto-interactions. In fact, the partition functions of an Hopfield model with auto-interactions and that of an Hopfield model without auto-interactions differ from sub-leading terms in N and so their thermodynamical properties have to be the same. Thus the presence of a new, unexplored, thermodynamical phase comprising multiple stable fixed points with non-vanishing basins of attraction at $\alpha \gg 1$ has to be ruled out.

Acknowledgement

Support from The Leverhulme Trust grant RPG-2013-48 is acknowledged. We wish to thank Pierfrancesco Urbani for insightful discussions.

Bibliography

[1] J J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

[2] J J Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the national academy of sciences*, 81(10):3088–3092, 1984.

[3] J J Hopfield, D I Feinstein, and R G Palmer. 'Unlearning' has a stabilizing effect in collective memories. *Nature*, 304:158 – 159, 1983.

[4] D O Hebb. *The organization of behavior: A neuropsychological theory*. Psychology Press, 2005.

[5] D J Amit, H Gutfreund, and H Sompolinsky. Spin-glass models of neural networks. *Physical Review A*, 32(2):1007, 1985.

[6] D J Amit, H Gutfreund, and H Sompolinsky. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Physical Review Letters*, 55(14):1530, 1985.

[7] Y Kabashima and D Saad. The TAP approach to intensive and extensive connectivity systems. *Advanced Mean Field Methods–Theory and Practice*, 6:65–84, 2001.

[8] M Mézard. Mean-field message-passing equations in the hopfield model and its generalizations. *Physical Review E*, 95(2):022117, 2017.

[9] V Folli, M Leonetti, and G Ruocco. On the maximum storage capacity of the hopfield model. *Frontiers in Computational Neuroscience*, 10:144, 2017.